arXiv:1403.7591v1 [cs.MM] 29 Mar 2014

Building A Large Concept Bank for Representing Events in Video

Yin Cui, Dong Liu, Jiawei Chen, Shih-Fu Chang
Department of Electrical Engineering, Columbia University

April 1, 2014

Abstract

Concept-based video representation has proven to be effective in complex event detection. However, existing methods either manually design concepts or directly adopt concept libraries not specifically designed for events. In this paper, we propose to build Concept Bank, the largest concept library consisting of $4,876$ concepts specifically designed to cover 631 real-world events. To construct the Concept Bank, we first gather a comprehensive event collection from WikiHow, a collaborative writing project that aims to build the world's largest manual for any possible How-To event. For each event, we then search Flickr and discover relevant concepts from the tags of the returned images. We train a Multiple Kernel Linear SVM for each discovered concept as a concept detector in Concept Bank. We organize the concepts into a five-layer tree structure, in which the higher-level nodes correspond to the event categories while the leaf nodes are the event-specific concepts discovered for each event. Based on such tree ontology, we develop a semantic matching method to select relevant concepts for each textual event query, and then apply the corresponding concept detectors to generate concept-based video representations. We use TRECVID Multimedia Event Detection 2013 and Columbia Consumer Video open source event definitions and videos as our test sets and show very promising results on two video event detection tasks: event modeling over concept space and zero-shot event retrieval. To the best of our knowledge, this is the largest concept library covering the largest number of real-world events.

1 Introduction

Representing and detecting complex events from unconstrained videos remains one of the most challenging problems in computer vision. By definition, an event is a complex activity that involves people interacting with other people and/or objects under certain scene settings. For instance, "changing a vehicle tire" can be defined as a complex event where

human, vehicle, tire and tools interact with each other in an outdoor environment. The existing research on video event detection focuses on the use of low-level features combined with sophisticated learning models, and achieves satisfactory performances to some extent [6, 17, 23]. However, these works fail to provide semantic information in a video event. This hampers high-level event understanding, especially when the number of training videos is small or zero. Therefore, a reasonable and computationally tractable way is to decompose videos depicting an event into a set of atomic concepts [14, 28]. These concepts can be perceived as building blocks of a complex event and are expected to provide a meaningful intermediate level of abstraction towards representing an event rather than low-level visual-audio features directly extracted from videos. In this paper, we focus on such an approach in which the textual title of the event (*e.g.*, "grooming an animal") is used as query to first find concepts relevant to the event, and then we use the selected concepts as representations in supervised event modeling or zero-shot learning.

There are two existing approaches to generate concept based video representation. The first is to manually define suitable concepts for each event, which involves too many human efforts and hence is not feasible for large-scale problems. The second is to directly utilize the existing banks built on certain ontologies. However, the main issue is that such banks are not specifically designed for complex events and do not contain enough relevant concepts. Blindly applying these banks with significant amount of irrelevant concepts will degrade the very purpose of high-level concept-based video representation.

This motivates us to build an event-oriented concept library called "Concept Bank" (CB) for high-level representation of complex events in video. In this paper, we identify three challenges for building such a concept library and make significant contributions in developing corresponding solutions summarized below.

The first challenge is that the library should be **quasi-comprehensive** and cover as many real-world events as possible. To address this challenge, we collect and select events from Wikihow [27], a collaborative writing project that aims to build the world's largest how-to manual for any possible event[1]. In this way, we end up with 631 events as our event collection. To the best of our knowledge, this is the largest event collection in the literature.

The second is how to **discover concepts** and build their visual models for each event. To solve this, we apply our recently proposed idea on automatic event-driven semantic concept discovery from Web images [4]. The web is a rich source of information with tremendous images captured for various events and these images are roughly annotated with descriptive tags that indicate the semantics of the image contents. Our intuition is that if we crawl

[1]Although WikiHow only focuses on how to do anything, we realize that it is created based on user's common interests and thus has good coverage on almost every aspect of human daily life. Advanced event collection techniques can be also used to extract visually detectable events from other knowledge bases, and easily incorporated into our current system.

enough number of images by an event query, tags of these images could somehow reveal certain semantics related to the event statistically, and suggest relevant concepts of the event. In light of this, we crawl images and their associated tags from Flickr using event queries extracted from WikiHow, then find the most relevant event-specific concepts from the tags and build visual model for each concept.

Finally, after we get a comprehensive concept library for all events, it remains unclear on how to choose the most **relevant concepts** to represent any possible query event. To address this task, we propose a semantic matching method to measure the relevance between each concept and the query event (without using any training videos). We then choose the top ranked concepts to represent the event and get a compact concept based video representation.

Figure 1 illustrates an overview of how we build the proposed Concept Bank and use it to represent an event video. Experimental results show that the video representation with only few hundred dimensions generated by CB achieve state-of-the-art event detection performance.

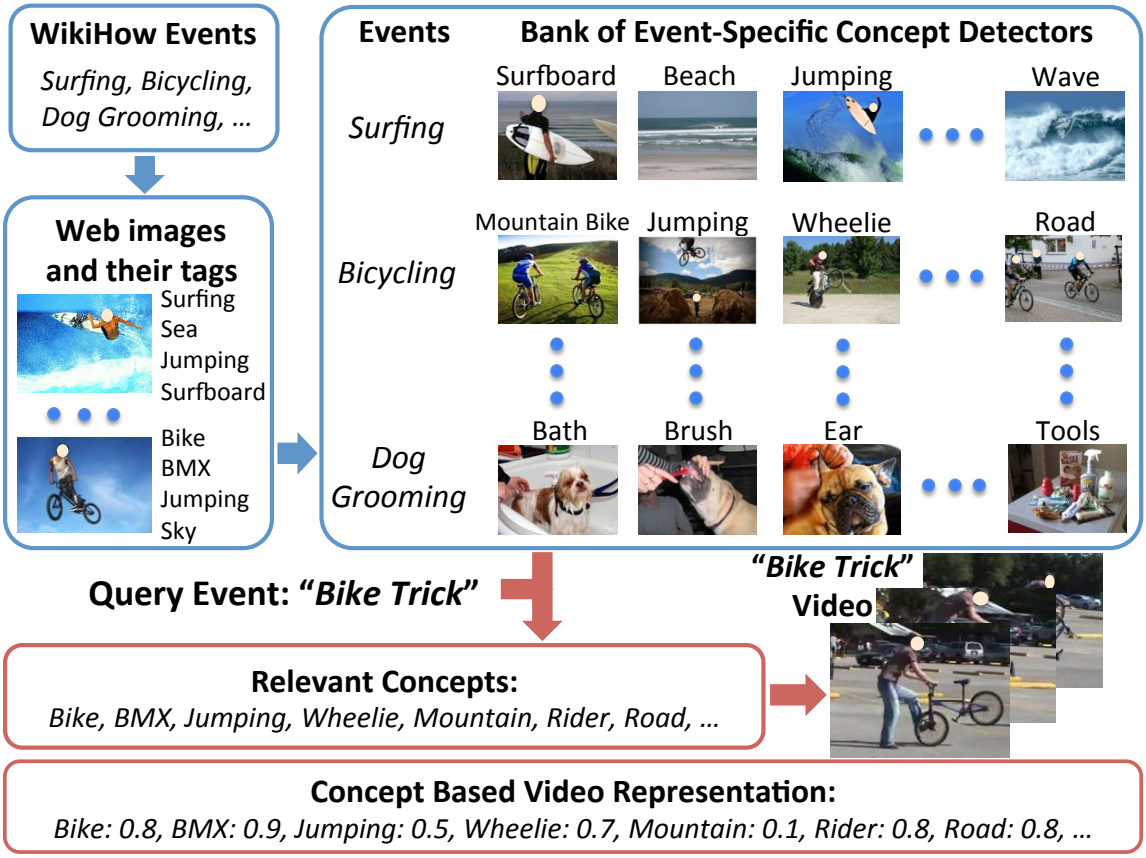

Figure 1: **Building Concept Bank** (in the upper three boxes): We collect a comprehensive event list from WikiHow. For each event, we search Flickr images and discover relevant concepts from their tags [4]. We train a concept detector for each discovered concept, resulting a large set of event-specific concept detectors. **Video Representation by Concept Bank** (in the lower part): Given a query event, relevant concepts are selected based on semantic matching (Section 5). Then the selected concept detectors are applied on frames of a video clip, generating a concept score vector on each frame. Finally, we use average-pooling to aggregate all concept score vectors across the frames.

2 Related Work

Video event detection has became an important research area in computer vision literature. Natarajan *et al.* [17] investigated the multimodal fusion of low-level video features and the extracted videotext information in video content, and achieved good performance on the detection. Ye *et al.* [29] discovered the bi-modal audio-visual codewords and leveraged the joint patterns across the audio and visual space to boost event detection performance. Duan *et al.* [6] incorporated the web sources videos crawled from Youtube to relieve the insufficiency of the number of training videos of an event, and developed a cross-domain video event detection model. Tang *et al.* [23] developed a large margin framework to exploit the latent temporal structure in successive clips of a long event video. These excellent works focus on modeling events into sophisticated statistical models or fusing mutimodal information. However, none of them can reveal the rich semantics in event videos.

There are some recent works that try to perform event detection with semantic concepts. Yang *et al.* [28] applied deep belief nets to group a large number of event video shots into a number of shot clusters, and then treat each cluster center as a data-driven concept. Then each video is mapped onto the cluster centers and encoded into a concept based representation. However, such data-driven concepts do not convey any semantic information and hence cannot be utilized for high-level semantic understanding. Liu *et al.* [14] manually defined concepts present in event videos and categorized them into "object", "scene" and "action". They annotated the presence of each concept on training videos and then built the individual concept detectors. Nevertheless, as aforementioned, this requires too many manual efforts, and is not applicable for real-world large-scale video event detection tasks. Different from these prior works, we focus on automatically discovering potential concepts present in any possible events with interpretable semantic meaning.

Notably, we applied our previous work on concept discovery from Internet images [4] as a building component of our concept library construction process. This previous work aims at discovering concepts from a pre-specified set of event definitions within a known domain and studying the large beneficial impacts of such event-specific concepts in event retrieval, zero-example detection, and summarization. In contrast, this paper focuses on discovering semantic concepts from an external event knowledge base, WikiHow, and using its rich ontological hierarchy to organize the large number of concepts learned from the Web. Such ontological structure plays an important role in handling novel events that have not been seen in the learning stage, as confirmed by the significant performance gains reported in the experiments.

There are some existing concept libraries built for different purposes. Object bank [12] consists of 200 objects taken from the intersection set of most frequent $1,000$ objects between image datasets LabelMe [21] and ImageNet [5]. Each object detector is trained with

$100 \sim 200$ images and their object bounding boxes. Classemes [24] is a concept library comprised of $2,659$ concepts defined from Large Scale Concept Ontology for Multimedia (LSCOM) [16]. Each concept detector is trained with 150 images from *bing.com* search engine using the LP-β multiple kernel learning algorithm. Action bank [22] is built for high-level representation of human activities in video, which contains 205 template actions. Although these libraries achieve good performances on different tasks, they are not designed for video events. In our work, we build a concept library specifically designed for video events, and the concept ontology and models are automatically discovered from the Web.

3 Building Concept Bank Ontology

In this section, we will introduce the procedure of building Concept Bank ontology, including event collecting and concept discovery for each event.

3.1 Collecting Events from WikiHow

In order to get a quasi-comprehensive list of events, we choose among the events from WikiHow [27]. WikiHow is a wiki, similar to Wikipedia, in which internet users can read or edit the articles. Currently, WikiHow contains 163, 957 articles that are organized into 20 categories in a hierarchical structure. The 20 categories cover the major aspects of human daily life including "Food and Entertaining", "Pets and Animals", "Sports and Fitness" and so on. Figure 2 illustrates a portion of the hierarchical structure of event category "Sports and Fitness" in WikiHow. Each category contains a number of subcategories (*e.g.*, "Individual Sports" and "Outdoor Recreation") that are instantiations of the higher level event catalog. Subcategories are further divided into various fine-grained events in hierarchy (*e.g.*, "Bicycling", "Mountain Biking" and "Fishing"), each has several articles (*e.g.*, "Jump a Mountain Bike" under "Mountain Biking"), which are not shown in the figure, describing the detailed procedure of each specific event.

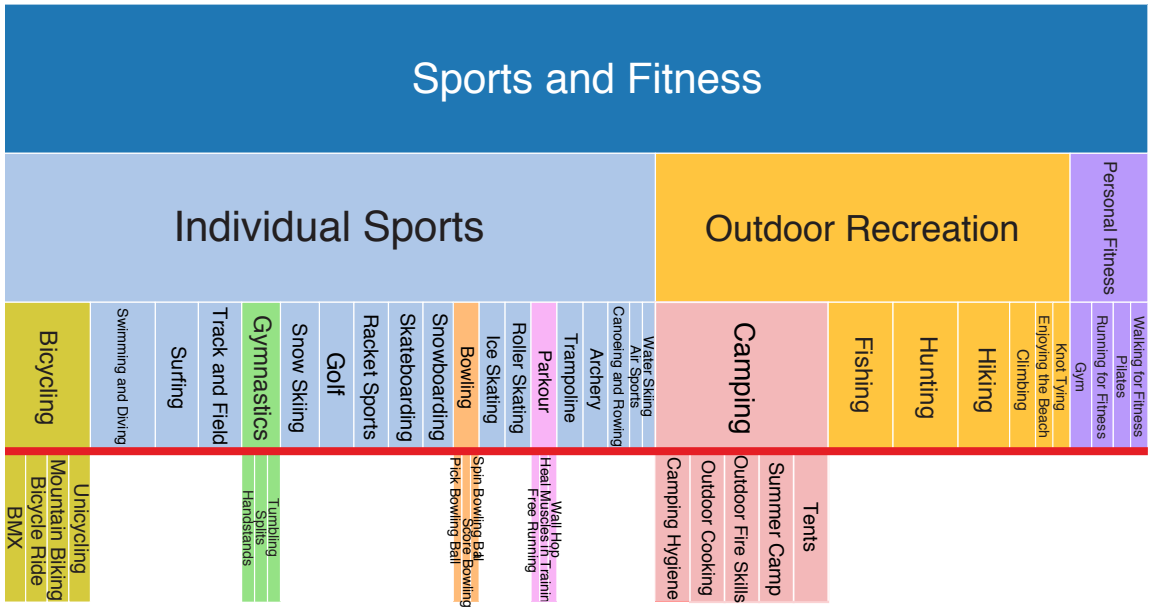

Figure 2: A portion of hierarchical structure for event category "Sports and Fitness" in WikiHow.

We find that some events in WikiHow are closely related and are sub-events within an event (*e.g.*, "Mountain Biking" and "Unicycling" under "Bicycling"). Our intuition is that organizing events into such fine granularity will result in a huge event collection with heavy semantic redundancy, which is impractical and unnecessary for our task. Therefore, we only focus on the first three layers in WikiHow (above the bold solid red line in Figure 2) and treat all events in the third layer as potential events in our library, resulting in $1,257$ event candidates in total.

Furthermore, we notice that some events are not visually detectable. For example, events "getting good grades in college" or "being successful" do not convey consistent visual patterns that can be modeled by computer vision techniques. Therefore, a filtering is performed to remove such undetectable events. To this end, we ask two human labelers to judge whether each of the $1,257$ events corresponds to a visually detectable event. After filtering, we end up with 631 events which are included as the events in our library.

3.2 Concept Discovery for Each Event

After collecting events from WikiHow, we need to associate a number of concepts to each event. One straightforward way is to use advanced Natural Language Processing technique to extract concepts from article contents on WikiHow. However, besides the complexity of the procedure itself, the resultant concepts do not correspond to any visual patterns, which also brings challenges for training data acquisition. On the other hand, the images on Flickr are associated with tags that describe the semantic of visual content in different events. Therefore, we leverage the tags of Flickr images to discover potential concepts for events.

In this work, we adopt our previous work on event-driven concept discovery from Internet images [4] to discover concepts for WikiHow events. For each of the 631 events, we use the article names under this event as textual queries to crawl images and their associated tags from *Flickr.com* image search engine. We only keep color images with resolution higher than 200×200 pixels. The crawling ends up with $727,910$ images in total and the average number of images is $1,154$ per event. Since some tags are not semantically meaningful words (such as a camera brand name) or do not have consistent visual patterns (such as "biology" or "economy"), directly accepting these tags as concepts may introduce lots of noises and degrade the performance. Therefore, we adopt the following two-phase process to discover concepts that have both semantic meanings and visually consistent patterns from tags.

Step I : Noisy Tag Cleansing. Notice that a tag could be either a word or a phrase. So firstly, we convert a tag into lower case and divide the tag into tokens by a tokenizer, then remove non-word tokens and tokens belonging to stop words (like "the", "and", or

"with") [3]. We then define a meaningless-word-list including camera names (like iPhone or Nikon) and people names (like Peter or Julia). For a tag, if one of its tokens appears in our meaningless-word-list or does not match any synonyms in WordNet, the tag is regarded as meaningless. During this step, we also bundle together tokens within a WordNet synonym with lemmatization [3]. For example, "run", "running", "runs" and "ran" will be regarded as a same token "run". By doing this, only the tags with semantic meaning are retained. Finally, we rank tags within each event in descending order by their frequencies and select top 100 tags as candidate concepts.

Step II : Concept Visualness Verification. If a concept is not visually related, the concept detector will be hard to generalize due to high intra-class variance [2]. Therefore, we use the cross-validation performance as the measure of visualness. Specifically, for a candidate concept, we collect images with the concept as positive samples and randomly select a same number of images which do not contain the concept as negative samples. Then we build a concept detector (see Section 4 for details) and calculate the 2-fold cross-validation performance. In this procedure, we adopt Average Precision (AP) as the performance evaluation metric. Finally, only candidate concepts with cross-validation APs higher than 0.8 are selected as concepts for the event.

In order to get statistically reliable concept detectors, we further do a post processing that removes those concepts contain less than 100 training images. By doing this, we end up with $4,876$ concepts. Note that the same concept may appear in different events (and are trained with different training images, see Section 4). Therefore, we call our discovered concepts as "event-specific concepts".

3.3 Concept Bank Ontology

Based on the collected events and concepts, we construct a five-layer tree structured concept ontology that helps event detection (see Section 5). In more details, the root layer is a single node covering all events and concepts in the Concept Bank. The second layer consists of 19 event categories ranging from "Arts and Entertainment" to "Travel", each corresponds to a category in WikiHow. In the third layer, we have in total 130 nodes in which each node corresponds to an event subcategory. In the fourth layer, we have 631 events. Furthermore, each of the 631 event nodes has several child nodes corresponding to its event-specific concepts. Therefore, the bottom layer of the CB tree ontology contains $4,876$ leaf nodes (concepts). Figure 3 illustrates a branch of hierarchy structure in CB ontology corresponding to event category "Sports and Fitness". Due to the space limitation, we remove the root node and only show at most 8 concepts for each event. In the next section, we will associate a concept detector with each of the leaf nodes.

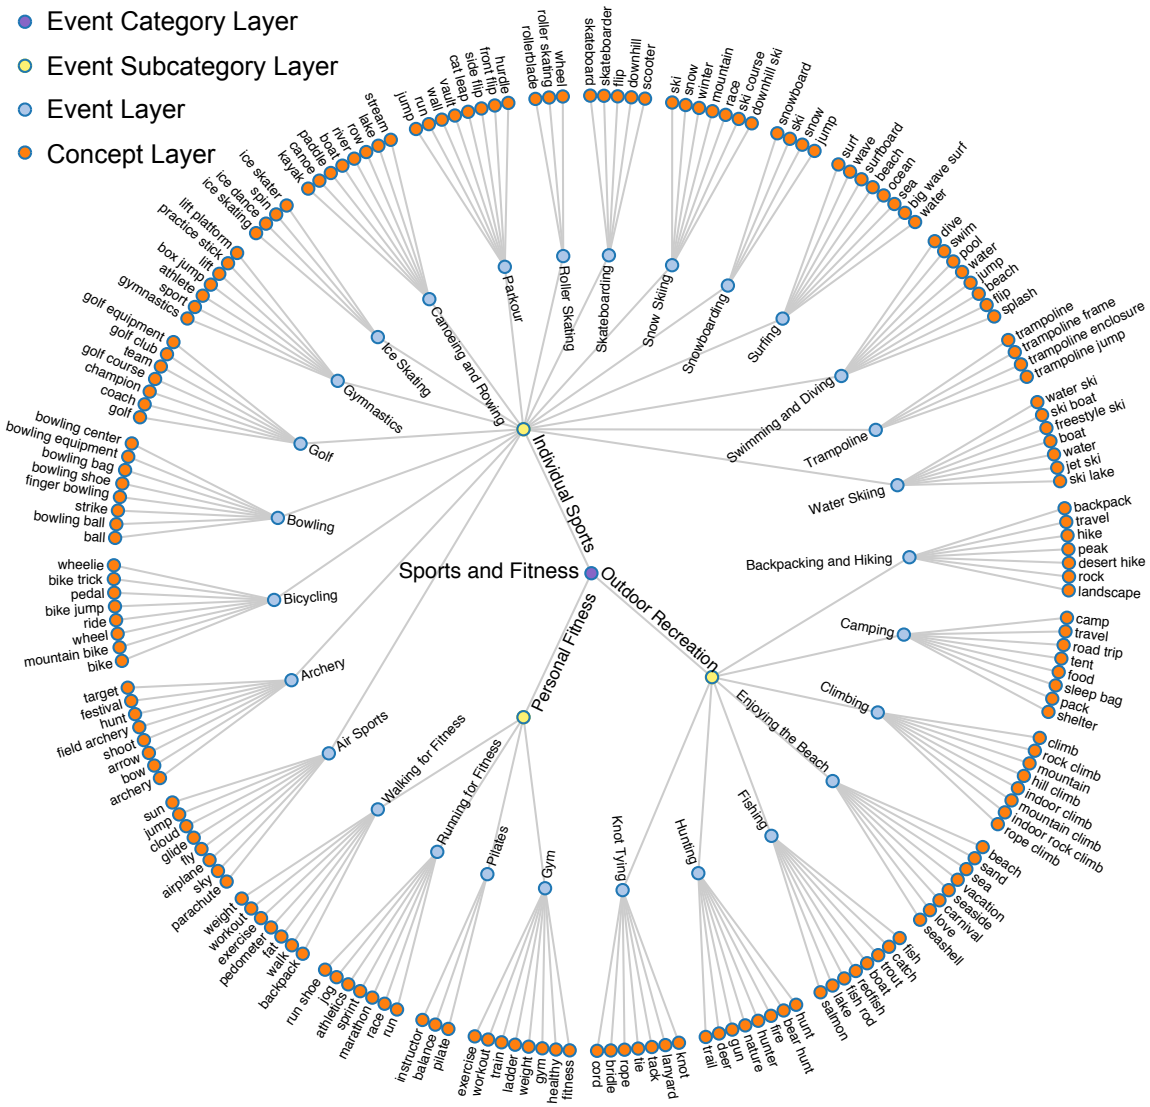

Figure 3: A branch of the hierarchy structure in Concept Bank ontology corresponding to event category "Sports and Fitness".

4 Learning Concept Detector

In this section, we will discuss how to build concept detector for each discovered concept.

4.1 Feature Extraction

In order to describe visual content information comprehensively and complementarily, we select 5 low-level feature descriptors that are proven to be effective in representing visual content. They are respectively **SIFT** [15], **GIST** [19], **Gabor** [8], **LBP** [18] and **Transformed Color Distribution** [25]. All descriptors are densely extracted on grids of 20×20 pixels with 50% overlap from images. For each type of the extracted descriptors, we train a codebook with $1,000$ codewords, and partition each image into 1×1, 2×2 and 3×1 blocks for spatial pyramid matching [10]. Finally, we adopt soft quantization [26] to represent each image as a $8,000$ dimensional histogram for each feature. We also adopt these features in generating concept scores for videos.

4.2 Training Image Selection

Since images and their tags are crawled from the web, images associated with a concept are noisy and contain some outliers. Therefore, we need to remove the noisy images before training concept detectors. Our intuition is that the majority of images in a concept form a common "visual theme" in the image search result, while distracting images tend to depart from such visual commonality. This motivates us to develop a collective approach to measure the relative confidence of each image in the search result. Specifically, for each concept c, we collect a set of images $F = \{1, \ldots, N\}$ that contain the concept as their tag, where N is the total number of images. For image i, we extract M features ($M = 5$ in this work) and denote its feature vector as $\mathbf{f}_m^i \in \mathbb{R}^{d_m}$ $(m = 1, \ldots, M)$ with d_m denoting the feature dimensionality of the m-th feature. We adopt Kernel Density Estimation (KDE) method [20] to estimate the probability of image i belonging to concept c defined as:

$$p(c|i) = \frac{1}{M \times N} \sum_{j=1}^{N} \sum_{m=1}^{M} \mathcal{G}_m(\mathbf{f}_m^i - \mathbf{f}_m^j), \tag{1}$$

where $\mathcal{G}_m(\cdot)$ is the Gaussian kernel function of the m-th feature defined as:

$$\mathcal{G}_m(\mathbf{f}_m^i - \mathbf{f}_m^j) = \exp(-\frac{\|\mathbf{f}_m^i - \mathbf{f}_m^j\|^2}{\sigma_m^2}), \tag{2}$$

in which $\|\cdot\|$ denotes the l_2 norm, σ_m is the kernel radius of the m-th feature setting as the mean value of all pairwise distance among the images.

We use $p(c|i)$ as the confidence score of image i belonging to concept c and select s images with highest confidence scores as the positive training samples for each concept. In this work, we set $s = 100$ and choose $t = 1,000$ images from other concepts as negative training samples. Figure 4 shows 5 top selected images for exemplary concepts, from which we can see that the selected images are highly reliable while at the same time achieve reasonable diversity.

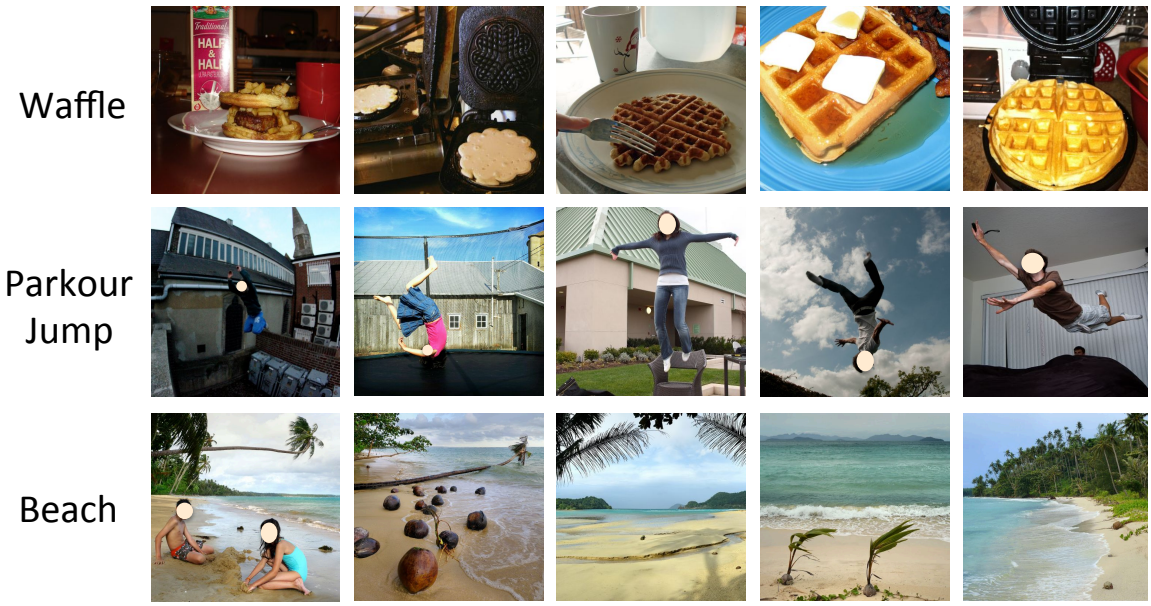

Figure 4: Top 5 selected Web images for 3 exemplary concepts using our training image selection method.

Since a concept may appear under different events, our method will select different training images for each event-specific concept detector. Therefore, each concept detector will convey context information for its belonging event.

4.3 Concept Detector Training

After selecting s positive and t negative training images for each concept, we use Multiple Kernel Linear SVM (MK-LSVM) classifier with L1-Loss to train our concept detector. The

dual norm of its objective function is:

$$\begin{aligned} \min_{\boldsymbol{\alpha},\boldsymbol{\beta}} \quad & \frac{1}{2}\boldsymbol{\alpha}^T\big(\sum_{m=1}^{M}\beta_m\mathcal{K}_m + \gamma I\big)\boldsymbol{\alpha} - \mathbb{1}^T\boldsymbol{\alpha} \\ \text{s.t.} \quad & \boldsymbol{\alpha}^T\boldsymbol{y} = 0,\ 0 \le \alpha_i \le C, \quad \forall i, \\ & \sum_{m=1}^{M}\beta_m = 1,\ \beta_m \ge 0, \quad \forall m, \end{aligned} \tag{3}$$

where $\mathbf{y} = [y_1, \ldots, y_{s+t}]^\top \in \{0,1\}^{s+t}$ is the label vector with $y_i = 1$ if the i-th image is positive and $y_i = 0$ otherwise, $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_{s+t}]^\top \in \mathbb{R}^{s+t}$ is the dual variable vector, $\boldsymbol{\beta} = [\beta_1, \ldots, \beta_M]^\top \in \mathbb{R}^M$ is the kernel weight vector, $\mathcal{K}_m \in \mathbb{R}^{(s+t)\times(s+t)}$ is the Gram matrix of the m-th feature, I is an identity matrix , $\mathbb{1}$ is the all-one vector, γ and C are model parameters. We use LIBLINEAR [7] and choose $\gamma = 0.01$ and $C = 1$ in training concept detectors.

5 Representing Videos with Concept Bank

Given a video clip, we can directly apply all $4,876$ concept detectors in Concept Bank to generate the concept representation. However, as aforementioned, it is desirable to choose the most relevant concepts related to the query event and generate a compact concept representation.

In our work, we turn to use the semantic similarity matching between event and concept names to find relevant concepts for an event query. In order to determine semantic similarity, we adopt ConceptNet [13], a semantic network developed by MIT Media Lab containing common sense knowledge for computers to capture human world. We use Divisi, a Python API for ConceptNet, to calculate semantic similarity between two words. The similarity between two phrases are the maximum similarity between all pairs of words in the two phrases.

Denote by $sim(a, b)$ the similarity between two words (phrases) a and b calculated from ConceptNet. For a given event e and a concept c_0, we take advantage of the hierarchy in the Concept Bank ontology to calculate their semantic similarity. Specifically, denote c_1, c_2 and c_3 as the three ancestors of concept c_0 in Concept Bank ontology residing at event layer, event subcategory layer and event category layer respectively. Then the semantic similarity $S(e, c_0)$ between e and c_0 can be calculated as follows:

$$S(e, c_0) = \prod_{i=0}^{3} sim(e, c_i). \tag{4}$$

For instance, for event "grooming a dog", we calculate its similarity with concept "brush" (similarity value 0.7) and its three ancestors: , "dog grooming" (similarity value 1), "dogs" (similarity value 1), and "pets and animals" (similarity value 0.8). Then the similarity between event "grooming a dog" and concept "brush" is $0.8 \times 1 \times 1 \times 0.7 = 0.56$.

Eq. (4) can be justified as follows. Since the same concept may appear under different events, involving high-level ancestors in the similarity calculation will measure the relevance of each concept based on the coherent similarity of its ancestors with respect to the event query, leading to more precise relevant concept selection. In this way, when answering event query "grooming a dog", we can effectively select concept "brush" from event "dog grooming" rather than from event "personal makeup" under "Personal Care and Style" even if it also contains concept "brush".

With semantic similarity matching, we choose the top n most relevant concepts for the query event and use their concept detectors to generate concept based video representation. Specifically, we evenly sample m frames from a video clip. Then we apply the selected concept detectors on each frame and adopt their probabilistic outputs as concept scores.

The final representation of the video is the average of concept score vectors across the sampled frames. We choose $n = 100$ and $m = 20$ in the experiments.

6 Experimental Evaluations

In this section, we evaluate the effectiveness of the Concept Bank in video event detection. We first introduce the datasets and comparison methods and then present experiment analysis on two event detection tasks. Specifically, we would like to evaluate whether the large number of concepts discovered using WikiHow events can be used to effectively detect events specified in some benchmark datasets. Additionally, we aim at comparing different ways of training the concept detectors and mapping events to concepts.

6.1 Dataset and Comparison Methods

Dataset. We evaluate our proposed CB representation on two video event detection datasets: (1) **TRECVID 2013 Multimedia Event Detection (MED) dataset**. This is a dataset of 32,744 videos over 20 event categories (i.e., the partition used in the Pre-Specified EK100 task [1]). These 20 event categories are (from E1 to E20 respectively): "*birthday party*", "*changing a vehicle tire*", "*flash mob gathering*", "*getting a vehicle unstuck*", "*grooming an animal*", "*making a sandwich*", "*parade*", "*parkour*", "*repairing an appliance*", "*working on a sewing project*", "*attempting a bike trick*", "*cleaning an appliance*", "*dog show*", "*giving directions to a location*", "*marriage proposal*", "*renovating a home*", "*rock climbing*", "*town hall meeting*", "*winning a race without a vehicle*", "*working on a metal crafts project*". We follow the standard partition of this dataset, which includes 7,787 videos in the training set and 24,957 videos in the test set. Achieving good performance on this dataset is quite challenging because the majority of videos in this dataset are background videos (only 2,000 videos in training set and 1,489 videos in test set belong to one of the 20 events). (2) **Columbia Consumer Video (CCV) dataset**. This is a dataset of 9,317 YouTube videos over 20 event categories, where 4,659 videos are used for training and 4,658 videos for testing [9]. These 20 categories are (from E1 to E20 respectively): "*basketball*", "*baseball*", "*soccer*", "*ice skating*", "*skiing*", "*swimming*", "*biking*", "*cat*", "*dog*", "*bird*", "*graduation*", "*birthday*", "*wedding reception*", "*wedding ceremony*", "*wedding dance*", "*music performance*", "*non-music performance*", "*parade*", "*beach*", "*playground*".

Comparison Methods. We compare the concept based video representations generated from the following methods. (1) **Classemes** [24], a 2,659 dimensional concept representation built on the LSCOM concept ontology [16]. We generate Classemes feature from each video frame and then average all frame features across a video as the video-level feature. (2) Selected Concepts from Best single Feature (**SCBF**). The concept detector is trained based on the best performing single feature, in which the training images are selected with the method described in Section 4.2. When a query event comes, it uses the semantic matching method in Section 5 to select relevant concepts for video representa-

tion. (3) Selected Concepts from Random Images (**SCRI**). The concept detector is trained by MK-LSVM, but the training images are randomly selected. The relevant concepts are selected using semantic matching. (4) Random Concepts from Selected Images (**RCSI**). The concept detector is trained by MK-LSVM with selected training images. However, concepts for the query event are randomly selected, instead of using semantic matching. (5) Our Selected Concepts from Selected Images (**SCSI**). The concept detector is trained by MK-LSVM with selected training images. The relevant concepts are selected based on semantic matching.

Evaluation Tasks. For each method, we evaluate the performance of different representations under two video event detection tasks: Event Detection in Concept Space and Zero-Shot Event Retrieval. In each task, the Average Precision (AP), which approximates the area under precision/recall curve, is used as evaluation metric on each event. We calculate mean Average Precision (mAP) over all event categories as the overall evaluation metric on the dataset.

6.2 Event Detection in Concept Space

In this task, all representations are regarded as high-level descriptors in concept space for video event detection. On each dataset, we choose 100 concepts for each of 20 events from the Concept Bank and concatenate the concept scores into a $2,000$-dimensional feature vector. A one-vs-all SVM classifier with χ^2 kernel is trained as the event detector. In the test stage, we use SVM probabilistic output as the event confidence score on each test video. Best parameters for SVM are chosen via 2-fold cross-validation on training set.

The comparison results between the five methods on MED and CCV are given in Figure 5, from which we have following observations: (1) The proposed SCSI method achieves the best performance by a significant margin compared with other methods, which demonstrates the effectiveness of SCSI in generating reliable concept based video representations. (2) The multiple feature based methods including SCSI and SCRI clearly beat the single feature based method SCBF. This shows the advantages of utilizing multiple features in concept detector training. (3) The SCSI significantly outperforms RCSI, which verifies that our semantic similarity matching method is able to select relevant concepts for event representation. (4) The SCSI has higher performance than SCRI. This is due to the fact that the former leverages more reliable training images in concept detector training than the random images utilized in the latter.

Figure 6 illustrates the variation of mAP with different number of concepts. For SCSI, SCRI, SCBF and RCSI, we choose the number of concepts from 10 to 100 with step of 10 for each event, and get concept numbers rang from 200 to $2,000$. For Classemes, we fix the concept dimension to be $2,659$ (lower dimension of classemes leads to worse

performance). As seen, the proposed SCSI method consistently outperforms all other methods when different numbers of concepts are adopted in each of the two datasets. This clearly demonstrates that our method is able to generate accurate concept scores on the videos to facilitate concept based event detection. Furthermore, the event detection performance keeps improving as the number of concepts increases. This shows that event modeling over concept space can benefit from the increasing number of relevant concepts. We also notice that, on both MED and CCV datasets, the proposed SCSI outperforms $2,659$ dimension Classemes even with only 200 dimensions (10 concepts for each of the 20 events), which again clearly demonstrates the advantages of our CB representation.

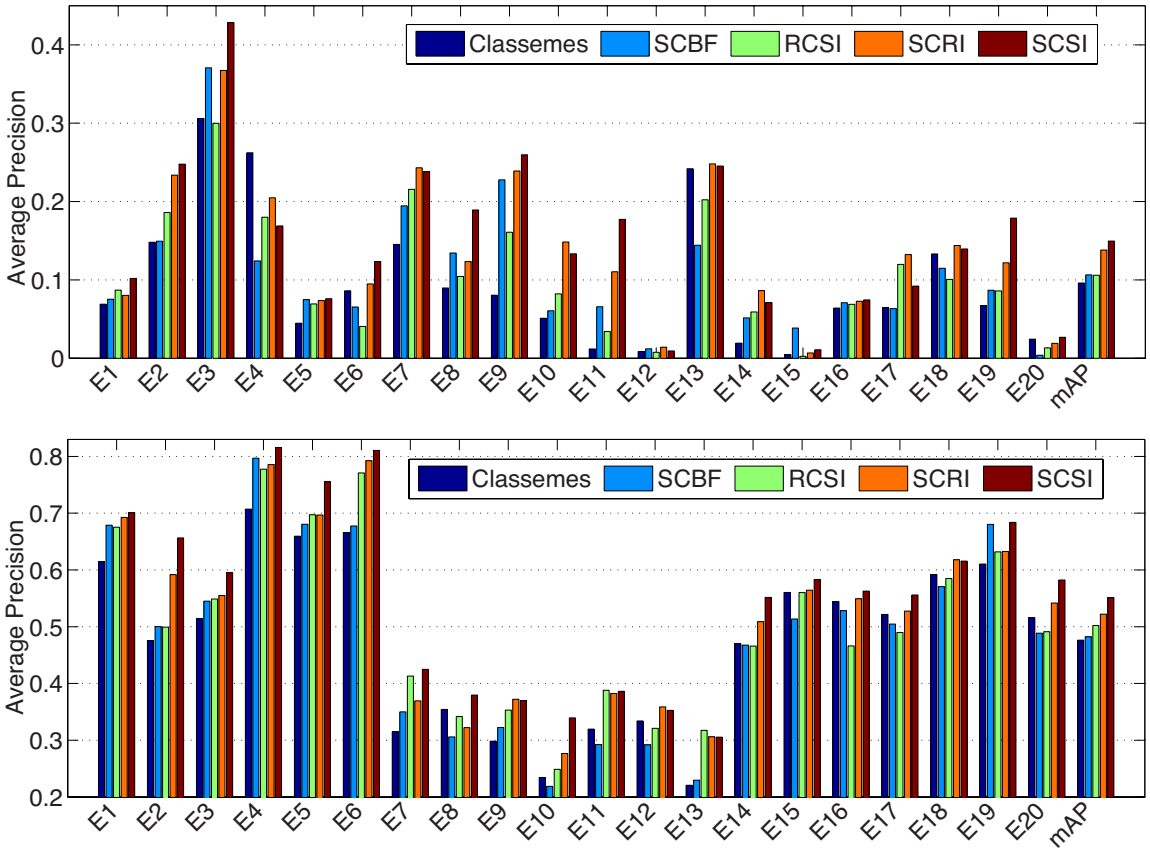

Figure 5: Performance comparison on event detection in concept space task (left: MED; right: CCV).

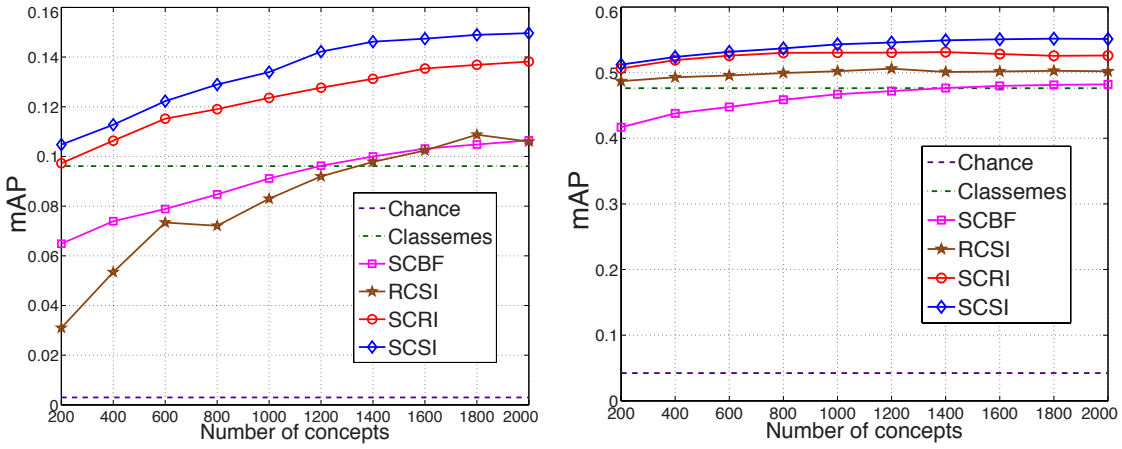

Figure 6: Event detection performance with different number of concepts (left: MED; right: CCV).

6.3 Zero-Shot Event Retrieval

In this task, we do not use any training videos, but directly use the concept scores on the test video to perform event retrieval. We call this zero-shot event retrieval since the procedure is purely semantic based without using any training examples. Specifically, we rank all test videos based on the detection scores of each selected relevant concept and then fuse the rank lists of all relevant concepts to generate the final rank list. In this work, we adopt the normalized rank fusion method [11] to combine multiple ranking lists. Specifically, for a test set with n videos, the fusion score $R(i)$ for i-th video is calculated by $R(i) = \frac{1}{d}\sum_{j=1}^{d}(1 - \frac{r_j}{n})$, where $r_j \in \{1, 2, \ldots, n\}$ is the rank position of i-th video in the rank list generated by j-th concept and d is the concept number. In this way, we get rid of the influence caused by numeric scale differences of raw concept scores generated from different concept detectors. Finally, we rank test videos based on the fusion scores.

In zero-shot retrieval, since the only available information is the query event name, we do not utilize concept names from other event categories. Therefore, we only choose 100 relevant concepts selected for the specific query event based on semantic similarity matching described in Section 5. To get a fair comparison, we also select 100 concepts from Classemes using the same method.

Figure 7 shows the comparison results between the five methods for zero-shot retrieval task and Figure 8 illustrates the variation of mAP with different number of concepts. There are two points we want to mention: (1) The performance of zero-shot event retrieval is worse than that of supervised event modeling over concept space (Section 6.2). This is because the latter uses training samples to obtain a more sophisticated event model while

the former is merely based on semantic score fusion. Notably, since the concepts are all randomly chosen and irrelevant to the query event, the performance of RCSI degrades a lot and is close to chance in this task. (2) Our SCSI method keeps improving as the number of concepts, whereas the performance of Classemes decreases. This indicates that Classemes does not have enough concepts related to the event, so the increasing number of irrelevant concepts will not help the detection. However, our Concept Bank has a broader coverage thus it can benefit from the increasing number of concepts.

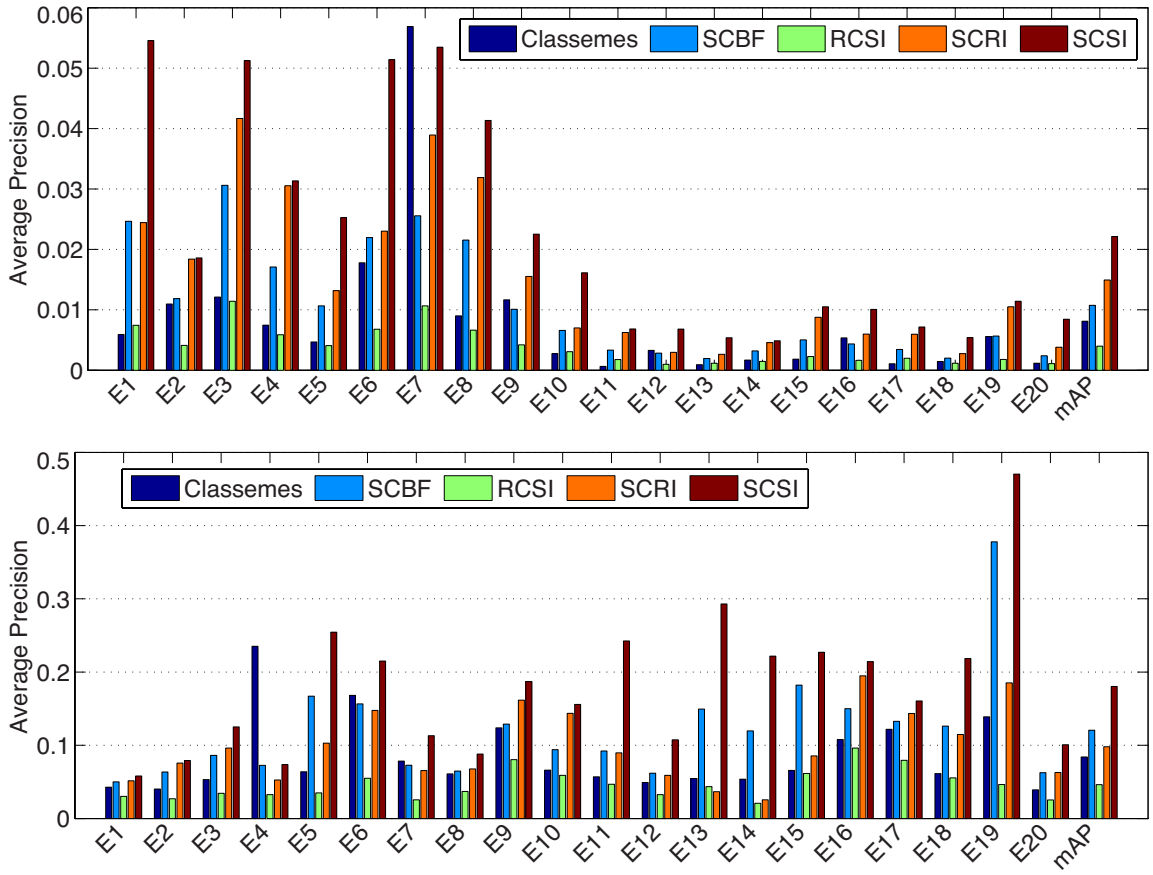

Figure 7: Performance comparison on zero-shot event retrieval task (left: MED; right: CCV).

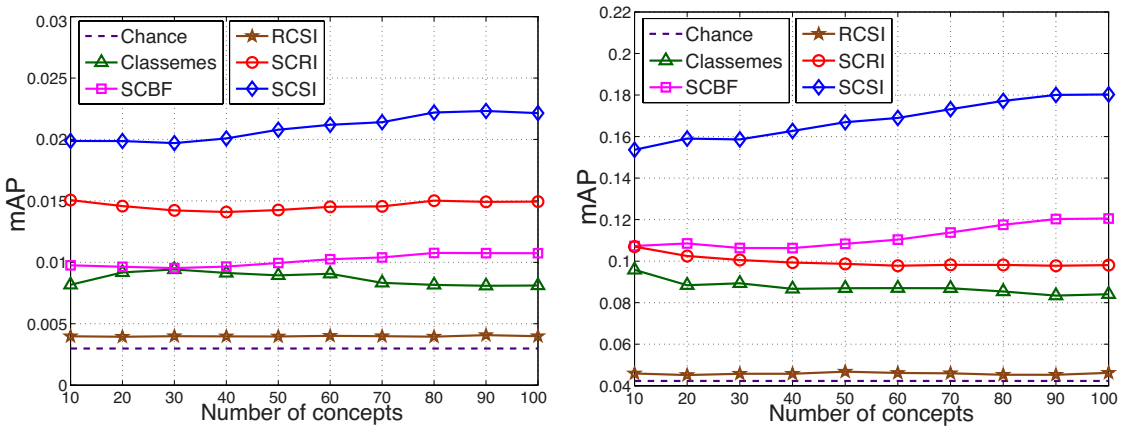

Figure 8: Zero-shot event retrieval performance with different number of concepts (left: MED; right: CCV).

6.4 Semantic Recounting in Videos

For each video of a target event, we rank all the concepts discovered for the event based on their confidence scores and treat the top ranked concepts as the semantic description of the video content. Such a procedure is able to reveal the semantic information contained in a video and is thus called video semantic recounting. Figure 9 shows the recounting result on videos from some exemplary events in MED and CCV, in which the top 5 ranked concepts generated by our method are selected as concepts for each video. As can be seen, these concepts reveal the semantics contained in the videos, which verifies the effectiveness of the our discovered concepts in representing video semantics.

6.5 Computational Cost

The proposed Concept Bank is efficient computationally. We calculate the running time on a 2.8GHz Windows workstation. In average, for a event query, semantic matching needs 78 seconds with Python. For a video, the average time to generate a 100 dimension concept score (including feature extraction) is around 110 seconds with MATLAB.

MED E5 Grooming An Animal: pet, cute, cat, dog, kitty

MED E6 Making A Sandwich: meal, dish, French toast, pan, recipe

MED E10 Working On A Sewing Project: yarn, crochet, knit, art, paint

CCV E2 Baseball: catch, tournament, baseball, competition, cheer

CCV E16 Musical Performance: concert, live music, musician, band, guitar

CCV E19 Beach: blue, island, beach, sand, sea

Figure 9: Event Video Recounting Results: each of the 6 rows shows evenly subsampled frames of an example video and the top 5 relevant concepts detected in the video.

7 Conclusions and Future Work

We have introduced Concept Bank, a concept library specifically designed for complex event representation in videos. The library consists of $4,876$ concepts organized in a five-layer tree structure, in which the higher-level nodes correspond to the event hierarchies and the leaf nodes are the event-specific concepts. To include as many real-world events in the library as possible, we collect all visually detectable events in WikiHow. For each event, we use its article names in WikiHow to search Flickr and then discover relevant concepts from the tags associated with the crawled images of this event. A Multiple Kennel Linear SVM classifier is then trained for each concept with the crawled images as the atomic concept detector. In addition, we also develop a semantic matching method to determine relevant concepts for a coming event query. Experiments over two event detection tasks verify the effectiveness of the proposed Concept Bank. For future work, we will further extend the Concept Bank with motion concepts, which are trained based on the spatial-temporal features extracted from videos.

References

[1] http://www.nist.gov/itl/iad/mig/med13.cfm. 16

[2] Tamara L Berg, Alexander C Berg, and Jonathan Shih. Automatic attribute discovery and characterization from noisy web data. In *ECCV*. 2010. 9

[3] Steven Bird. Nltk: the natural language toolkit. In *COLING/ACL on Interactive presentation sessions*, 2006. 9

[4] Jiawei Chen, Yin Cui, Guangnan Ye, Dong Liu, , and Shih-Fu Chang. Event-driven semantic concept discovery by exploiting weakly-tagged internet images. In *ICMR*, 2014. 2, 4, 5, 8

[5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5

[6] Lixin Duan, Ivor W Tsang, Dong Xu, and Stephen J Maybank. Domain transfer svm for video concept detection. In *CVPR*, 2009. 2, 5

[7] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. Liblinear: A library for large linear classification. *JMLR*, 2008. 13

[8] David J Field et al. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A*, 1987. 11

[9] Yu-Gang Jiang, Guangnan Ye, Shih-Fu Chang, Daniel Ellis, and Alexander C Loui. Consumer video understanding: A benchmark database and an evaluation of human and machine performance. In *ICMR*, 2011. 16

[10] Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, 2006. 11

[11] Joon Ho Lee. Analyses of multiple evidence combination. In *ACM SIGIR*, 1997. 19

[12] Li-Jia Li, Hao Su, Li Fei-Fei, and Eric P Xing. Object bank: A high-level image representation for scene classification & semantic feature sparsification. In *NIPS*, 2010. 5

[13] Hugo Liu and Push Singh. Conceptneta practical commonsense reasoning tool-kit. *BT technology journal*, 2004. 14

[14] Jingen Liu, Qian Yu, Omar Javed, Saad Ali, Amir Tamrakar, Ajay Divakaran, Hui Cheng, and Harpreet S Sawhney. Video event recognition using concept attributes. In *WACV*, 2013. 2, 5

[15] David G Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004. 11

[16] Milind Naphade, John R Smith, Jelena Tesic, Shih-Fu Chang, Winston Hsu, Lyndon Kennedy, Alexander Hauptmann, and Jon Curtis. Large-scale concept ontology for multimedia. *Multimedia, IEEE*, 2006. 6, 16

[17] Pradeep Natarajan, Shuang Wu, Shiv Vitaladevuni, Xiaodan Zhuang, Stavros Tsakalidis, Unsang Park, and Rohit Prasad. Multimodal feature fusion for robust event detection in web videos. In *CVPR*, 2012. 2, 5

[18] Timo Ojala, Matti Pietikainen, and Topi Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *TPAMI*, 2002. 11

[19] Aude Oliva and Antonio Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *IJCV*, 2001. 11

[20] Emanuel Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 1962. 11

[21] Bryan C Russell, Antonio Torralba, Kevin P Murphy, and William T Freeman. Labelme: a database and web-based tool for image annotation. *IJCV*, 2008. 5

[22] Sreemanananth Sadanand and Jason J Corso. Action bank: A high-level representation of activity in video. In *CVPR*, 2012. 6

[23] Kevin Tang, Li Fei-Fei, and Daphne Koller. Learning latent temporal structure for complex event detection. In *CVPR*, 2012. 2, 5

[24] Lorenzo Torresani, Martin Szummer, and Andrew Fitzgibbon. Efficient object category recognition using classemes. In *ECCV*. 2010. 5, 16

[25] Koen EA Van De Sande, Theo Gevers, and Cees GM Snoek. Evaluating color descriptors for object and scene recognition. *TPAMI*, 2010. 11

[26] Jan C van Gemert, Cor J Veenman, Arnold WM Smeulders, and J-M Geusebroek. Visual word ambiguity. *TPAMI*, 2010. 11

[27] By Wikipedians. *All About wikiHow*. PediaPress. 2, 7

[28] Yang Yang and Mubarak Shah. Complex events detection using data-driven concepts. In *ECCV*. 2012. 2, 5

[29] Guangnan Ye, I-hong Jhuo, Dong Liu, Yugang Jiang, D. T. Lee, and Shih-Fu Chang. Joint audio-visual bi-modal codewords for video event detection. In *ICMR*, 2012. 5